




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04L 12/437, 12/56		A1	(11) International Publication Number: WO 00/13376
			(43) International Publication Date: 9 March 2000 (09.03.00)
(21) International Application Number: PCT/US99/19488 (22) International Filing Date: 26 August 1999 (26.08.99) (30) Priority Data: 09/143,464 28 August 1998 (28.08.98) US (71) Applicant: INTEGRAL ACCESS, INC. [US/US]; 321 Billerica Road, Chelmsford, MA 01824 (US). (72) Inventor: NELSON, Gary, A.; 20708 North Deerpath Road, Barrington, IL 60010 (US). (74) Agent: HEFFAN, Ira, V.; Testa, Hurwitz & Thibault, LLP, High Street Tower, 125 High Street, Boston, MA 02110 (US).		(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: REDUNDANT PATH DATA COMMUNICATION



(57) Abstract

A redundant packet data communication system includes a transmitter transmitting a first packet and a second packet. The first packet has a first label indicating a receiver and a first path. The first packet also has a first packet identifier and payload. The second packet has a second label indicating the receiver and a second path. The second packet has a second packet identifier and payload substantially similar to said first packet's identifier and payload. The system also includes a receiver that receives the first packet via the first path. The receiver receives the second packet via the second path. The receiver determining from the first packet identifier and the second packet identifier that the packet payloads are substantially similar, and discards one of the packets.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

REDUNDANT PATH DATA COMMUNICATION

Priority Information

This application claims priority to U.S. Patent Application Serial No.: 09/143,464 filed on August 28, 1998.

Technical Field

5 This invention relates generally to the field of telecommunications and, more particularly, to the providing of increased reliability in telecommunications networks.

Background Information

In the field of telecommunications, telecommunications service providers typically provide a single dedicated circuit between two endpoints. Some of these dedicated circuits are
10 used to carry voice traffic and some dedicated circuits are used to carry data traffic. Telecommunications service providers have also offered what is referred to as 1+1 redundant service, which is the provision of two or more circuits between the same endpoints. One of the circuits is used for communication, and if that circuit fails, another circuit is used to communicate. Such allocation of circuits is useful because it provides an alternative path for the
15 telecommunications traffic, and insures that a dedicated circuit will be available to carry the traffic. Such allocation is inefficient and expensive, however, since one of the circuits is always inactive. In some implementations, the redundant circuits are allocated along different physical wire paths that are strung or laid along different physical routes. In this way a physical failure on one wire, for example due to an accidental wire cut, only affects one dedicated circuit, and does
20 not affect the remaining redundant circuits.

A telecommunications service provider traditionally offers such redundant service which both the active circuit and the redundant circuits have equal bandwidth. For example, a service provider might offer 1+1 service for a T-1 customer by providing two or more T-1 service circuits. This is inefficient and expensive if the customer does not use the full capacity of the
25 service circuit or if the customer requires 1+1 redundant capability for only for a portion of the customer's traffic that is significantly less than the capability of the service circuit.

The present invention addresses the need for redundancy in telecommunications circuits while reducing the inefficiencies associated with the use of such circuits.

- 2 -

Summary of the Invention

A method and system for providing redundant communications uses a packet-switched network to provide redundant service. The method and system can provide redundant service only for the traffic that requires redundant service, without the bandwidth waste associated with an unused dedicated circuit. The method and system provides the additional capability of automatic use of a functioning circuit when one of the communications paths fails, so that the traffic is unaffected by the failure.

In general, in one aspect, the invention features a method for redundant packet data communication. The method includes transmitting a first packet and a second packet. The first packet has a first packet identifier, first packet data, and a first label indicating a receiver and a first path. The second packet has a second packet identifier and payload identical to the first packet identifier and payload. The second packet also has a second label indicating the receiver and a second path. The first packet is received via the first path and the second packet is received via the second path. The method includes determining from the first packet identifier and the second packet identifier that the packet payloads are identical, and discarding one of the packets.

In general, in another aspect, the invention features a system for redundant packet data communication. The system includes a transmitter for transmitting a first packet and a second packet. The first packet has a first packet identifier and payload, first packet data, and a first label indicating a receiver and a first path. The second packet has a second packet identifier and payload identical to the first packet identifier and payload. The second packet also has a second label indicating the receiver and a second path. The system also includes a receiver for receiving the first packet via said first path and the second packet via said second path. The receiver determines from the first packet identifier and the second packet identifier that the first packet payload and second packet payload are identical, and discards one of the packets having identical identifier and payload.

In general, in another aspect, the invention features a method for communicating over a network. The method includes adding an identifier to a data packet, communicating a copy of the packet to a destination via a first network route, communicating a copy of the packet to said destination via a second network route, and using, at the destination, the copy of the of the packet that arrives first at the destination.

- 3 -

In general, in another aspect, the invention features a ring of network nodes for redundantly communicating label-switched data packets that have a label and data. The ring includes at least three nodes, with each node in communication with two other nodes such that the communication path formed by the at least three nodes is a ring. Each node is connected to each adjacent node by a communications link for transmitting label-switched data packets from the node to the adjacent node. Each node is also in communication with each adjacent node with a communications link for receiving label-switched data packets from the adjacent node. In the ring, a first one of the nodes sends two label-switched data packets with identical data to a second one of the nodes substantially simultaneously in two different directions around the ring in response to the labels in the label-switched data packets.

Embodiments of this aspect of the invention include the following features. In one embodiment, the second node uses the first of each packet with identical data received and discards the second. In another embodiment, the second node preferentially uses packets received from one direction around the ring. In another embodiment, the label-switched data packets comprise a plurality of micropackets.

The foregoing and other objects, aspects, features, and advantages of the invention will become more apparent from the following description and from the claims.

Brief Description of the Drawings

In the drawings, like reference characters generally refer to the same parts throughout the different views. Also, the drawings are not necessarily to scale, emphasis instead generally being placed upon illustrating the principles of the invention.

FIG. 1 is a block diagram of an embodiment of a packet-switched network capable of providing redundant service constructed according to the invention;

FIG. 2 is a block diagram of an embodiment of a unidirectional packet-switched ring capable of providing redundant service constructed according to the invention;

FIG. 3 is a block diagram of the MLPPP extensions to the PPP protocol utilized in an embodiment of the invention;

FIG. 4 is a block diagram of the encapsulation of a data packet in an embodiment of the invention;

FIG. 5 is a flowchart of procedure followed by a transmitting node in an embodiment of the invention;

- 4 -

FIG. 6 is a flowchart of procedure followed by a receiving node in an embodiment of the invention;

FIG. 7 is a flowchart of the procedure followed by a primary flow receiver in an embodiment of the invention; and

5 FIG. 8 is a flowchart of the procedure followed by a protection flow receiver in an embodiment of the invention.

Description

The invention relates to redundant service that is provided by use of a packet switched communications network. Both voice and data are sent redundantly by different paths in the packet switched network. The use of a packet switched network allows the redundant packets to
10 share the network bandwidth with other, non-redundant packets. If, as in one embodiment, the packet switched network is a label switched network, paths through the network are specified by the association of a label with each path.

Referring to FIG. 1, a packet-switched network 2 includes network nodes A-J. In one
15 embodiment, each node A-J is a label switch node, and the network is a label switched network. In a label switched network, network nodes make forwarding decisions based on a label associated with each packet. The label associated with each packet indicates that the packet is a member of a particular forwarding equivalence class ("FEC"), which is the set (or "class") of packets that are treated the same way by a network node, regardless of the packets' ultimate
20 network destination. A packet with one label is forwarded to the destination associated with a particular FEC, and a packet with another label is forwarded to the destination associated with another FEC, which may be the same or a different destination. The labels may be "swapped" by a label switch node, meaning that the label of an incoming packet may be changed before the packet is forwarded. Examples of technologies that use label switching include, but are not
25 limited to, such technologies as Multiprotocol Label Switching ("MPLS") as described in the Internet Engineering Task Force's Network Working Group Internet Drafts, Cell Switching Router ("CSR") technology as developed by Toshiba, IP Switching as developed by Ipsilon, Tag Switching as developed by Cisco Systems, and Aggregate Route-based IP Switching ("ARIS") as developed by International Business Machines Corporation.

30 The implementation of physical links or connections between the nodes is not a limitation on the scope of the invention. In various embodiments, the connections between the nodes

- 5 -

include, but are not limited to such communications technologies as voice-band modems, RS-232 serial, xDSL, ISDN, Ethernet, Firewire, ATM, DS-1/E-1, GR-303, and SONET/SDH connections. In some embodiments, nodes within the same network are connected with different types of physical connections.

5 To demonstrate the invention, an example of redundant service is described using the network of FIG. 1. Redundant service is provided by sending duplicate packets from node A to node D substantially simultaneously. For each packet to be transmitted from node A to node D, one of the duplicate packets is transmitted from node A to node D via the path of nodes A-B-C-D. Another of the duplicate packets is transmitted from node A to node D via the path of nodes
10 A-G-F-E-D. If both paths are functional, node D will receive two copies of the same packet, one via path A-B-C-D, and the other via path A-G-F-E-D. The duplicate packets contain an identifier so that node D can use one of the duplicate packets and ignore or discard the other duplicate packet. Because the line is not dedicated, the packet-switched connections between the nodes can carry other traffic besides the redundant traffic sent between nodes A and D. If node
15 A requires redundant service for some data, and does not require redundant transmission for other data, only the data that requires redundant service will be sent in duplicate. The other data can be sent via path A-B-C-D, path A-G-F-E-D, or even some other path.

In one embodiment, node D receives duplicate packets. The first (in time) duplicate packet that is received is used, and other duplicate packets are discarded. Thus, in reconstructing
20 the data, some data packets may be used that are communicated via one path, while other packets may be communicated via another, redundant route. In this embodiment, the packets that arrive first are used. If there is a failure along one of the paths, the traffic will continue to be communicated uninterrupted, because the packets from the other path will continue to arrive. Those packets from the other, operational path will be considered the first to arrive, and so they
25 will be used.

In another embodiment, node D waits until both duplicates are received before using one of the packets. Node D compares the packets to verify that they are identical before using one of the packets. In this way, node D verifies the integrity of the packets. If one of the packets is not received by a certain deadline or timeout, then just the first packet received is used.

30 In one embodiment, node D will use the first of each of the duplicate packets that is received, and will also monitor whether all duplicates have been received, to determine if there is

- 6 -

a network problem along any of the paths. If a network problem is detected, various actions may be taken. Again, this action is transparent to the users, and the traffic is unaffected, because a redundant path is still operational even while action is taken to correct or work around a network failure. Using the earlier example of communication between nodes A and D, if the path
5 between nodes E and F fails, redundant network users will not be affected since the redundant path A-B-C-D is still operational.

In one embodiment, a network problem is detected when a number of duplicate packets are not received, and a system manager is alerted to the problem. In another embodiment, upon detecting a network problem, node D will attempt to correct the problem by requesting that the
10 failed redundant communication be continued via another path. Depending on the configuration of the network, this may be accomplished by sending an alert or a request to the source, which is Node A in this example, or by sending an alert or request to another node.

Although the above example describes two redundant paths, the number of redundant paths is not a limitation of the scope of the invention. The description above can be extended to
15 include additional duplicate packets transmitted over additional duplicate paths. For example, in one embodiment a third duplicate is sent over a third path, and in one such embodiment, the first duplicate received of the three is used. As another example, in another embodiment, ten duplicates of each packet are sent over ten paths.

Referring to FIG. 2, in one embodiment, a ring of nodes W, X, Y, Z are each connected to
20 their neighbors in the ring by two unidirectional links. For example, node W has a unidirectional connection WX to node X by which node W can transmit data to node X, and node W has a unidirectional connection XW from node X by which node W can receive data from node X. In one embodiment, the connections from each node in the ring to its two neighbors are by a different wire route, so that a cable cut will not cut off communication with both neighbors at the
25 same time. For example, in one embodiment, the cable(s) carrying connection WX and XW are strung or laid in a different physical place or route than the cable(s) carrying connections ZW and WZ.

The unidirectional connections can be of various types and configuration. In one embodiment, each unidirectional connection is a Synchronous Optical Network ("SONET")
30 connection. For example, in one embodiment, each connection is the connection protocol PPP running over a SONET OC-3 connection on fiberoptic cable. In another embodiment, the

- 7 -

connection is a PPP connection running over a Synchronous Digital Hierarchy ("SDH") connection over fiberoptic cable at rate STM-1 (155Mb/s). While these embodiments are specific examples, other connection protocols, connections and physical links are within the spirit and scope of the invention.

5 In one embodiment, the nodes are label switches capable of receiving labeled packets and passing the labeled packets on to the next switch in the ring. The nodes are also capable of directing the packets out of the ring at each node. For example, node W is capable of passing packets out of the ring to node w'. Each node W-Z can transmit packets to each other node in either direction. For example, packets from network w' to network y' can be
10 transmitted from node W to node Y by the path W-WX-X-XY-Y. Node W can also transmit packets to node Y by the path W-WZ-Z-ZY-Y.

 In one embodiment, data to be transmitted with redundant service from one node to another is divided up into packets, an identifier is associated with each packet, and duplicate copies of each packet are sent at approximately the same time in both directions around the ring,
15 so that a copy of each packet is sent via each of the two paths around the ring from the source node to the destination node. For example, in the embodiment of FIG. 2, packets from node W to node Y are sent both by the path W-WX-X-XY-Y and the path W-WZ-Z-ZY-Y. As another example, packets sent from node Z to node Y are sent by path Z-ZY-Y and by path Z-ZW-W-WX-X-XY-Y.

20 As described with regard to the embodiment of FIG. 1, the receiving node will receive two copies of each packet. The identifier associated with each packet indicates to the receiver which of the packets received are duplicates of each other. In one embodiment, the receiving node uses the first packet received, and discards or ignores the second copy of the packet when it is received. In other embodiments, the duplicate packet is used.

25 In the ring embodiment, it is not necessary for all data from one node to another to be transmitted redundantly. Only the data for which 1+1 redundancy is required need be sent redundantly. This saves bandwidth and allows other data to be sent along with the redundant data. For example, less important data can be sent without redundancy between the nodes.

 In one embodiment, an identifier is included in each data packet. The identifier allows
30 the receiver to track whether duplicate copies are received for each packet. The identifier enables the receiver to use a packet and to ignore or discard the duplicate. In one embodiment, the

- 8 -

identifier is included in the link layer protocol. In other embodiments, other identifiers are associated with packets.

Referring to FIG. 3, in one embodiment, Point to Point Protocol ("PPP") is used as the link protocol between nodes, and the PPP Multilink Protocol ("MP") extensions are used to include identifiers in the PPP packets. PPP is a link layer protocol that includes a header 106 including address information 110, control information 112, and a packet type indicator 114. The Multilink PPP ("MLPPP") extensions include a four byte header 102, 104 which includes a twenty-four bit sequence number. This sequence number is used to identify the packets. As described in the Internet Engineering Task Force Request for Comments ("RFC") No. 1990, which specifies Multilink PPP, the sequence numbers are intended to be used to order packet fragments that are transmitted over multiple channels. In the context of redundant transmission, the MLPPP headers are used as individual packet identifiers, so that duplicate packets can be matched.

In one embodiment, redundant communications are implemented on a label switched network, such as a network that supports multiprotocol label switching ("MPLS"). In Multiprotocol Label Switching networks, labels are associated with data heading to the same destination. Forwarding, which is the passing of packets from node to node, is simplified by use of short fixed length labels to identify the forwarding equivalence class. Forwarding may require simple functions such as looking up a label in a table, swapping labels, and possibly decrementing and checking a time to live counter, but is much less complicated than routing of the sort that occurs in Internet Protocol routers. This is because the path is set up once with the assignment of labels. In addition to simplified forwarding, MPLS also provides for efficient explicit switching. With such explicit switching, the source node sets up a path through the network. This is accomplished through the use of a label distribution protocol in which the nodes communicate their positions in the network, and adjacent nodes agree to forward packets with a particular label. Once the connection has been set up, the act of transmitting a packet with the appropriate label sends the packet to its destination along that specified path. The paths can be set up in a number of ways, depending on the label-switching implementation. Each path is created by the creation of a forwarding equivalence class ("FEC") for that path. Insofar as a forwarding decision is concerned, all packets that get mapped into the same FEC are indistinguishable. Generally, a label distribution protocol is used to map labels to paths. A

- 9 -

label-switched node uses the label distribution protocol to inform other nodes of the bindings of each label to a particular FEC. MPLS architecture thus allows a node to request from its next hop a label binding for a particular forwarding equivalence class.

In one embodiment, the transmitter initiates the creation of at least two paths to a destination, a first path and a second path, through the network. Each path has an associated label. Transmission of a packet with one label will send a packet to the destination via one path, and transmission of a packet with another label will send a packet to the destination via another path. In one embodiment, the paths travel along completely different physical wire routes, so that a wire cut or other network error between nodes will not affect both paths. In another embodiment, there is some overlap in the physical wire routes of the paths.

Referring to FIG. 4, the data that is to be transmitted redundantly can be any sort of data. Typically, the data to be transmitted already will be encapsulated in higher level protocol packets. For example, the data can be encapsulated in high-level protocols, such as TCP/IP or other ISO Layer 3 and above protocols, but that is not a requirement. The data in the data packet 100 is not relevant to the scope of the invention, any data can be transmitted redundantly.

In one embodiment, the data is divided into PPP packets 110 by an MLPPP protocol stack such that each packet has a PPP header 106, MLPPP information 104, a PPP Checksum 108, and an associated sequence number 102. In the embodiment of FIG. 4, two copies 120, 121 of each MLPPP packet are transmitted, one with a first label 125 specifying the first path, and one with a second label 126 specifying the second path. The packets 120, 121 are transmitted from label-switched node to label-switched node until they reach their destination. The receiver can identify which packets came from which path, since the packets have different labels. The receiver can also determine which packets are duplicates, since the duplicates have the same MLPPP sequence number.

Referring to FIG. 5, in one embodiment, a network node receives data to be transmitted redundantly from a source (STEP 150). In one embodiment, the source is a module within the node. In other embodiments, the source is another computer or node in communication with the network node. In one embodiment, the data is in the form of packets. In other embodiments, the data is in various other forms, and requires division into packets. The data is divided up into packets, if necessary, and an identifier is attached to each packet (STEP 151). The identifier allows the receiving node to determine which packets are duplicates. For example, in one

- 10 -

embodiment, the identifier is a packet number. The packets, with identifier, are then inserted as duplicate packets into packets having different path information (STEP 152). In one embodiment, each packet is encapsulated in a label-switched packet having a different label. Each different label directs the packet along a different path. Finally, the packets are transmitted
5 (STEP 153). In one embodiment, the label-switched packets with different labels will be transmitted along different paths.

Referring to FIG. 6, in one embodiment, a receiver receives data in the form of a packet from a redundant source (STEP 160), for example from the source of FIG. 5. The receiver extracts the identifier from the packet (STEP 161). The receiver determines whether the packet
10 is a duplicate of a packet already received (STEP 162). In one embodiment, the receiver makes this determination by comparing the identifier to a list of identifiers already received. In one embodiment, if the packet is a duplicate the receiver records the arrival of the duplicate packet (STEP 163). This step is useful only to the extent that the receiving node tracks the communications performance of the various paths. For example, in one embodiment the receiver
15 records only the arrival of the duplicate packet. In another embodiment, the receiver records information indicating that the duplicate packet arrived, and how long the duplicate packet arrived after the first packet. In other embodiments, other information about the duplicate packet is recorded. In this context, recording can include, but is not limited to storing the information in a list or database and/or transmitting the information to another node for reporting to a system
20 operator, compilation of statistics, or storage in a list or database on another system. . Once information about the duplicate packet has been recorded, the duplicate packet is discarded (STEP 164). Alternatively, in one embodiment, the duplicate packet is stored.

As described above, the step of recording information about the duplicate packet, such as whether it was received and when it was received, is useful for tracking the status and
25 performance of the redundant communications paths. In an embodiment in which the receiver does not track the communications capability of the paths, the recording step, STEP 163, is not performed.

If the packet is not a duplicate of a packet already received (STEP 162), but is in fact the first packet received with the packet's identifier, the packet will be used (STEP 165). In one
30 embodiment, the packet is sent to a different module. In another embodiment, the packet is stored until all the packets in a set are received, and then the data is recombined. Also, in one

- 11 -

embodiment, information and statistics about the receipt of the packet may be recorded, for comparison to the duplicate packets and to determine when all the packets in a set have been received.

Referring to FIG. 7, in another embodiment, one path is determined to be the primary path, and another path is determined to be a protection path. In this embodiment, a receiver will preferentially use data from one path, and will use data from another path only to "fill in" packets that are missing. If the error rate associated with data transmission on the primary path reaches a predetermined threshold, the receiver will use the protection path as the primary path, and report the error to a system operator.

Still referring to FIG. 7, the receiver receives packets from the primary path (STEP 170) and extracts the identifiers (STEP 171). The receiver determines if a packet is missing from the data transmitted over the primary path (STEP 172). If all packets have been received, the duplicate packets received over the protection path are discarded (STEP 173). If a packet is missing, and the duplicate of that packet is received over the protection flow, then the receiver uses the packet from the protection flow (STEP 174). The error is reported or logged (STEP 175). An error rate counter is incremented (STEP 176), and if the error rate is greater than a predetermined threshold (STEP 177), then the receiver will switch to using a protection path as the primary path (STEP 178). The path switch can also cause a notification or logging of the event.

Referring to FIG. 8, the receiver receives the data transmitted over the protection path (STEP 190) and extracts the identifiers from the packets (STEP 191). The receiver will determine if packets are missing from the data transmitted over the protection path (STEP 192), and if all packets have been received, discard the duplicates (STEP 193). If packets are missing, the receiver will report or log the error condition (STEP 194), and may also track the frequency of the errors for reporting or logging.

Variations, modifications, and other implementations of what is described herein will occur to those of ordinary skill in the art without departing from the spirit and the scope of the invention as claimed. Accordingly, the invention is to be defined not by the preceding illustrative description but instead by the spirit and scope of the following claims.

What is claimed is:

- 12 -

Claims

- 1 1. A method for redundant packet data communication, comprising the steps of:
2 transmitting a first packet and a second packet, said first packet comprising a first packet
3 identifier, first packet data, and a first label indicating a receiver and a first path, said second
4 packet comprising a second packet identifier and payload identical to said first packet's identifier
5 and payload, and a second label indicating said receiver and a second path;
6 receiving said first packet via said first path;
7 receiving said second packet via said second path;
8 determining from the first packet identifier and the second packet identifier that the
9 packet payloads are identical; and
10 discarding one of the packets.
- 1 2. A system for redundant packet data communication, comprising:
2 a transmitter for transmitting a first packet and a second packet, said first packet
3 comprising a first packet identifier and payload, and a first label indicating a receiver and a first
4 path, said second packet comprising a second packet identifier and payload identical to said first
5 packet identifier and payload, said second packet further comprising a second label indicating
6 said receiver and a second path;
7 a receiver for receiving said first packet via said first path and said second packet via said
8 second path, determining from the first packet identifier and the second packet identifier that the
9 first packet payload and second packet payload are identical, and discarding one of the packets
10 having identical identifier and payload.
- 1 3. A method for communicating over a network, comprising the steps of:
2 adding an identifier to a data packet;
3 communicating a copy of the packet to a destination via a first network route;
4 communicating a copy of the packet to said destination via a second network route;
5 using, at said destination, the copy of the of the packet that arrives first at said
6 destination.
- 1 4. A ring of network nodes for redundantly communicating label-switched data packets,
2 said data packets comprising a label and data, said ring comprising:
3 at least three nodes, each node in communication with two other nodes such that the
4 communication path formed by the at least three nodes is a ring,

- 13 -

5 wherein each node is connected to each adjacent node by a communications link for
6 transmitting label-switched data packets from the node to the adjacent node and wherein each
7 node is in communication with each adjacent node with a communications link for receiving
8 label-switched data packets from the adjacent node; and

9 wherein a first one of the nodes sends two label-switched data packets with identical
10 data to a second one of the nodes substantially simultaneously in two different directions
11 around the ring in response to the labels in the label-switched data packets.

1 5. The ring of claim 4 wherein the second node uses the first of each packet with identical
2 data received and discards the second.

1 6. The ring of claim 4 wherein the second node preferentially uses packets received from
2 one direction around the ring.

1 7. The ring of claim 6 wherein the label-switched data packets comprise Multiprotocol
2 Label Switched (MPLS) data packets.

1 8. The ring of claim 6 wherein the second node preferentially uses packets received from
2 one direction around the ring until more than a predetermined number of errors occur, and then
3 the second node uses packets received from the other direction around the ring.

1/7




FIG. 1




FIG. 2




FIG. 3


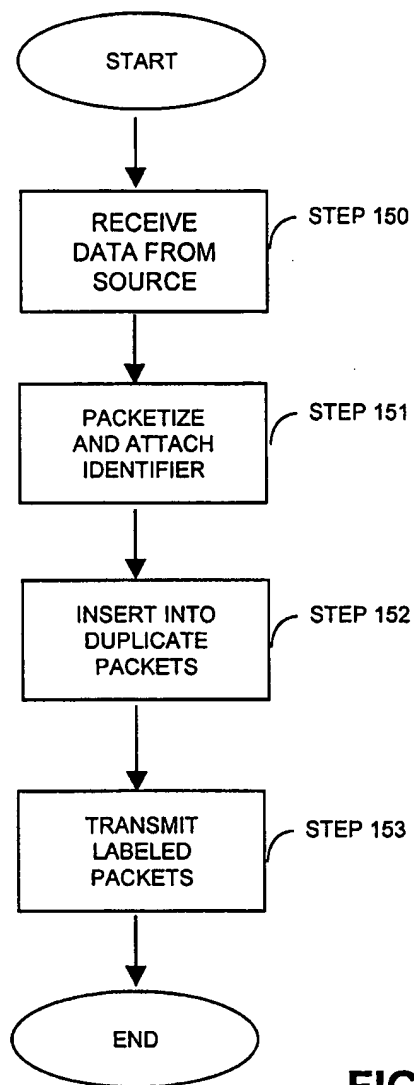
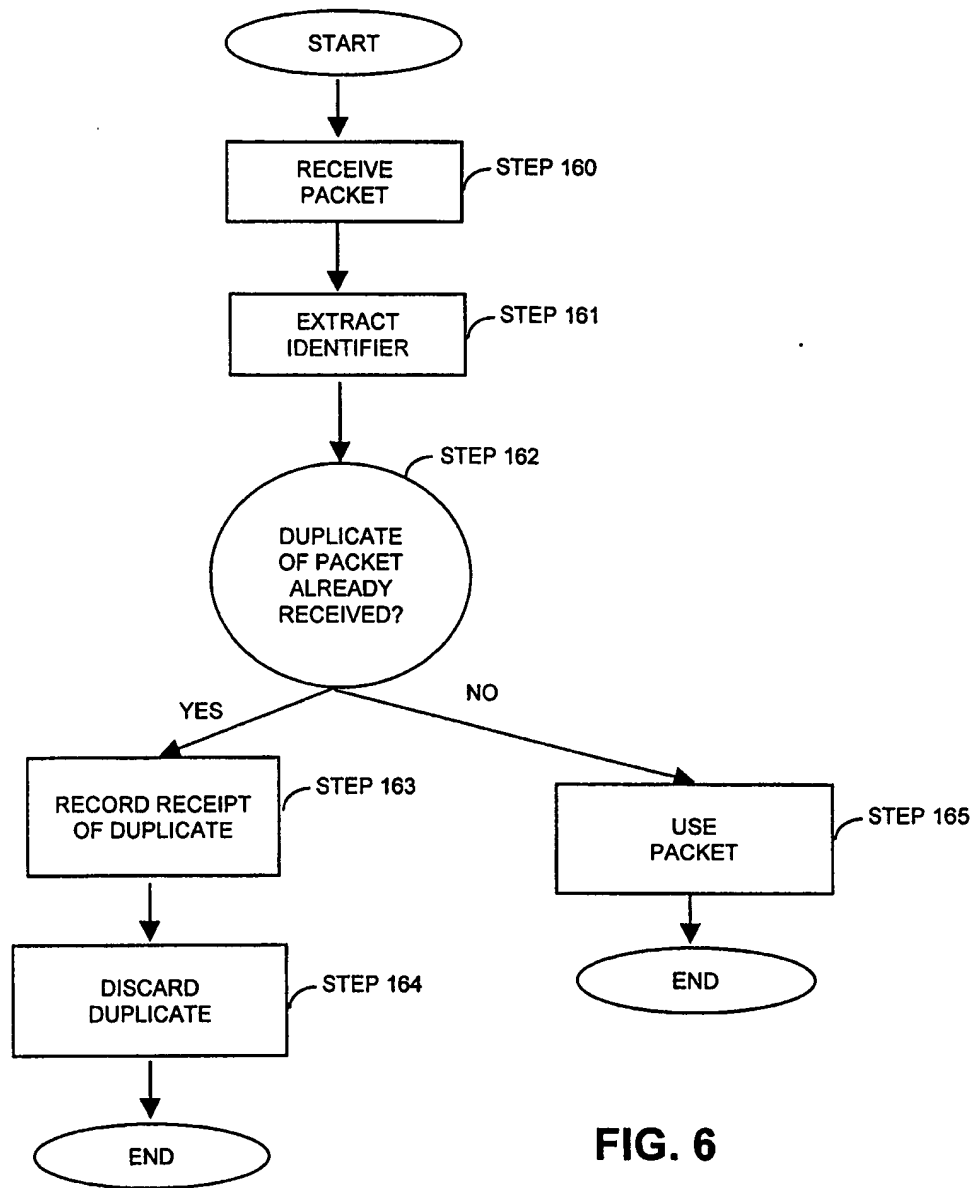


FIG. 4

**FIG. 5**

**FIG. 6**

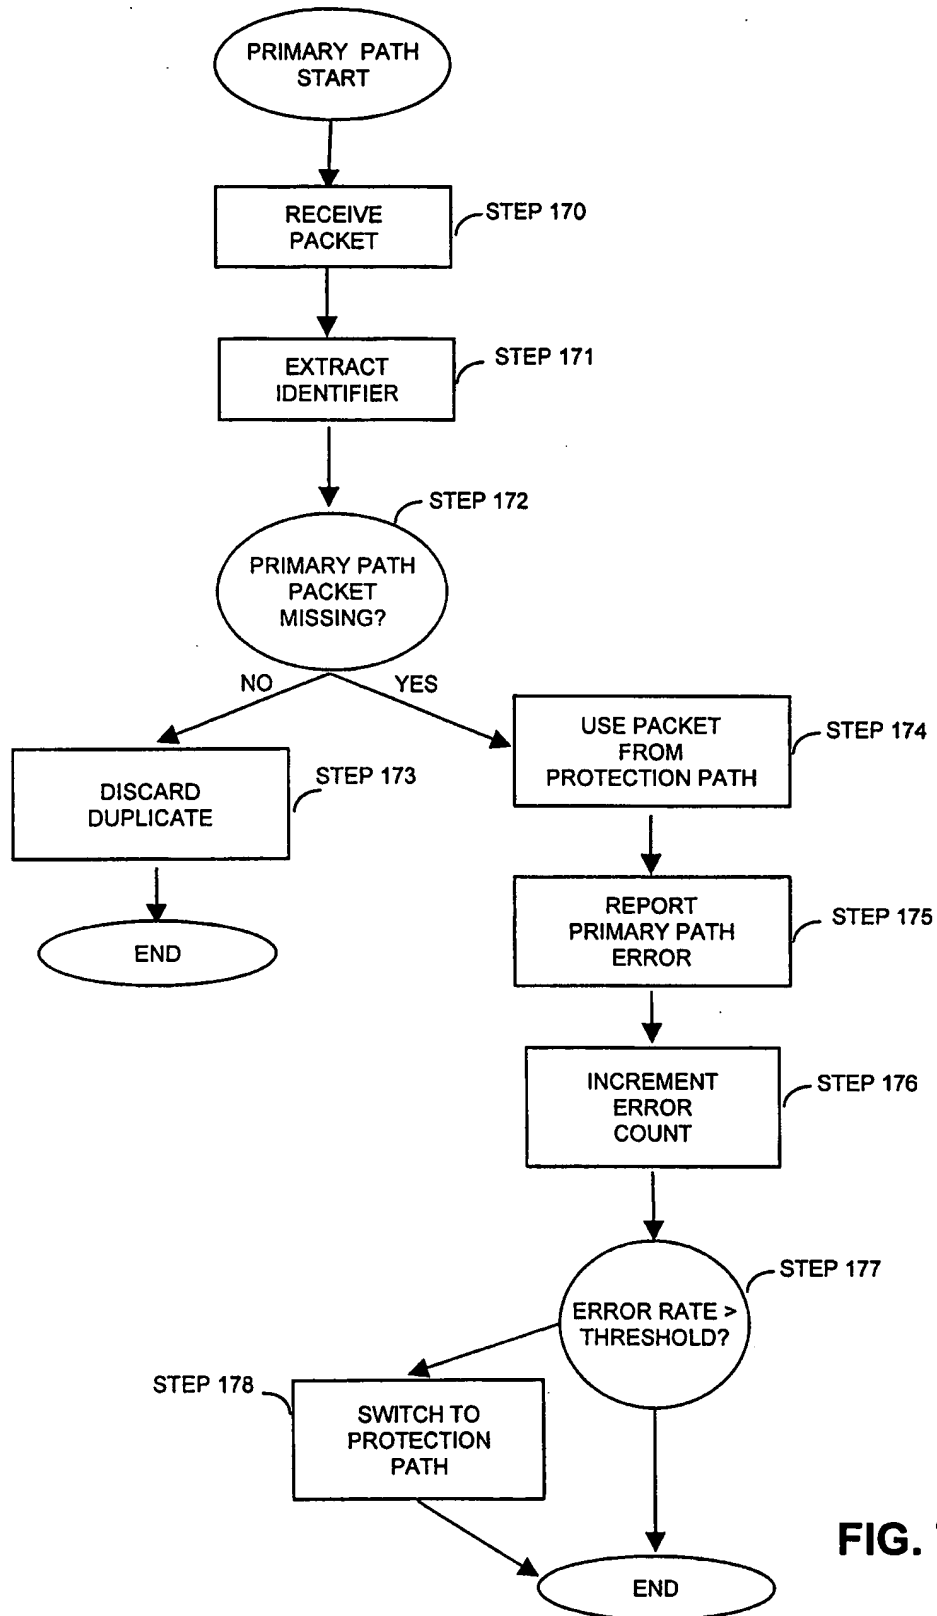


FIG. 7

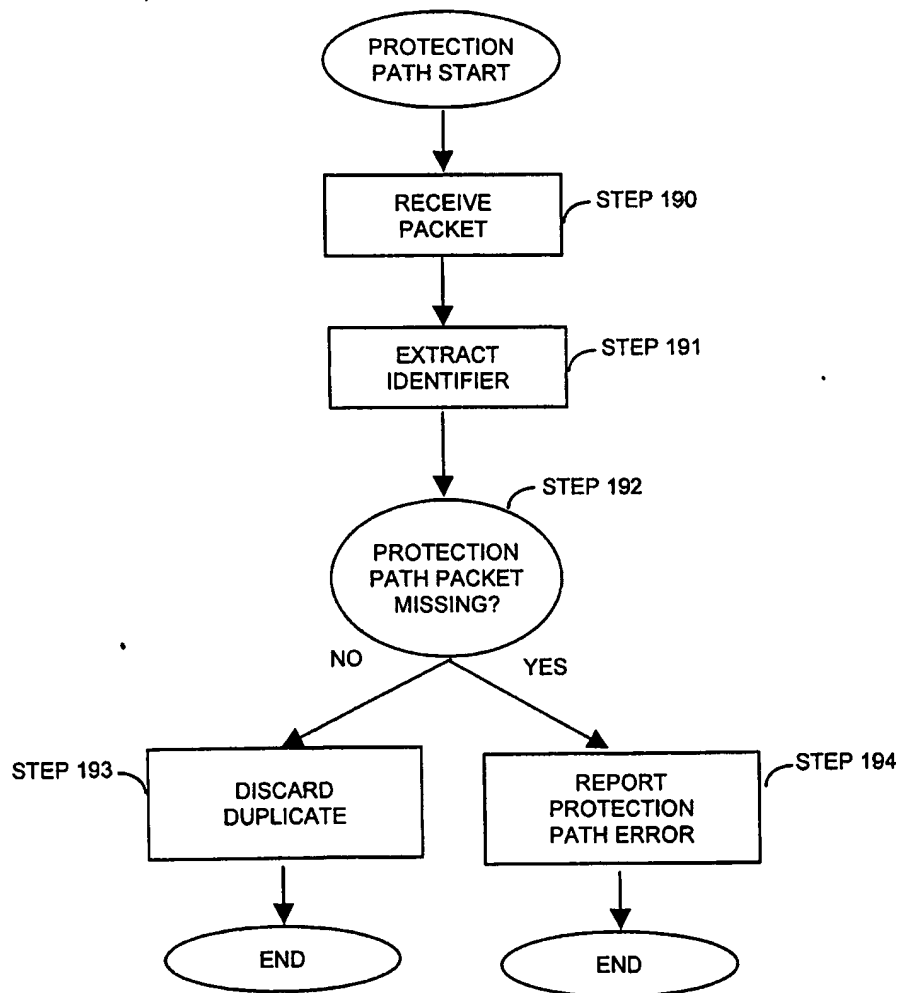


FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/19488

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/437 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>ORDA A ET AL: "ROUTING WITH PACKET DUPLICATION AND ELIMINATION IN COMPUTER NETWORKS" IEEE TRANSACTIONS ON COMMUNICATIONS,US,IEEE INC. NEW YORK, vol. 36, no. 7, 1 July 1988 (1988-07-01), pages 860-866, XP000570702 ISSN: 0090-6778 abstract page 860, left-hand column, paragraph 3 -right-hand column, paragraph 3 page 862, right-hand column, paragraph 1 --- -/--</p>	3

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

2 February 2000

Date of mailing of the international search report

10/02/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040. Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Blanco Cardona, P

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/19488

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>KAO B ET AL: "AGGRESSIVE TRANSMISSION OF SHORT MESSAGES OVER REDUNDANT PATHS" IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS,US,IEEE INC, NEW YORK, vol. 5, no. 1, 1 January 1994 (1994-01-01), pages 102-109, XP000433594 ISSN: 1045-9219 abstract page 102, right-hand column, paragraph 1 -page 103, left-hand column, paragraph 6</p>	1-3
A	<p>US 5 559 959 A (FOGLAR ANDREAS) 24 September 1996 (1996-09-24) column 3, line 34 -column 4, line 9 column 6, line 4 - line 8 column 7, line 41 - line 46</p>	1,2
A	<p>SY K -B K ET AL: "SOURCE ROUTING FOR LOCAL AREA NETWORKS" PROCEEDINGS OF THE GLOBAL TELECOMMUNICATIONS CONFERENCE AND EXHIBITION(GLOBECOM),US,NEW YORK, IEEE, vol. -, 1985, pages 3411-3415, XP000619287 abstract page 1022, left-hand column, paragraph 7 - paragraph 8</p>	1,2
Y	<p>US 5 187 709 A (WILLIAMSON GREGORY L ET AL) 16 February 1993 (1993-02-16) abstract column 1, line 10 - line 15</p>	4
A	<p>column 2, line 9 - line 27 column 3, line 58 - line 68 column 8, line 52 - line 65</p>	5-8
Y	<p>EP 0 468 813 A (NIPPON ELECTRIC CO) 29 January 1992 (1992-01-29) abstract column 1, line 10 - line 24</p>	4
A	<p>claim 1</p>	7

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/19488

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5559959 A	24-09-1996	DE 4331579 A CA 2132068 A EP 0645919 A JP 7170277 A	30-03-1995 17-03-1995 29-03-1995 04-07-1995
US 5187709 A	16-02-1993	AU 6287890 A WO 9117609 A	27-11-1991 14-11-1991
EP 0468813 A	29-01-1992	JP 2663687 B JP 4084535 A CA 2047949 A,C DE 69114203 D DE 69114203 T US 5150356 A	15-10-1997 17-03-1992 28-01-1992 07-12-1995 04-04-1996 22-09-1992